

IN THE CLAIMS

What is claimed is:

1. A service operating on a computer system to process requests from client processes, comprising:

5 a primary service instance operating on the computer system, including means for receiving and processing client requests, from the client processes;

at least one backup service instance operating on the computer system, including means for receiving and processing the client requests, wherein each backup service instance is logically equivalent to the primary service instance;

10 wherein the primary service instance includes

means for determining which external requests, when processed by the secondary service instance, are logical requests such that processing of the determined logical requests cause the external view of each backup service instance to change; and

15 means for communicating with each backup service instance to provide an indication of the determined logical requests to each backup service instance.

2. The service of claim 1, wherein each backup service instance has a physical behavior different from the physical behavior of the primary service instance.

20 3. The service of claim 1, wherein a failure domain of the primary service instance is substantially independent of a failure domain of each backup service instance.

4. The service of claim 1, wherein the determined logical requests providing means of

the primary service instance provides the determined logical requests to each backup service instance via a network protocol.

5. The service of claim 1, wherein each backup service instance includes means for persistently storing the determined logical requests.

6. The service of claim 1, wherein each backup service instance includes means for communicating, to the primary service instance, acknowledgement of the determined logical requests.

7. The service of claim 6, wherein the primary service instance includes means for committing the determined logical requests in response to the acknowledgements of the determined logical requests.

8. The service of claim 7, wherein each backup service instance includes means for committing the results of processing the determined logical requests.

9. The service of claim 8, further comprising:
order ensuring means for ensuring that, for a sequence of determined logical requests, each backup service instance commits the results of processing the determined logical requests in the same order as the primary service instance commits the results.

10. The service of claim 9, wherein:

the order ensuring means includes means for maintaining a table of determined logical requests; and

each backup service instance includes means for processing the table of determined logical requests to determine an order of committing results of processing the determined logical requests.

11. The service of claim 10, wherein:

the means for maintaining a table of determined logical requests includes means for associating sequence values with the determined logical requests; and

the means for processing the table includes means for processing the sequence values.

12. The service of claim 1, wherein:

the primary service instance includes means for receiving an indication of a path failure between the primary service instance and a client process that occurs after a result of processing a particular client request is committed and before the client process receives the result from the service; and

the primary service instance includes means for receiving the particular client request again and for providing the response to the client request without again processing the client request.

13. The service of claim 1, further comprising:

each backup service instance includes means for receiving an indication of a failure of the primary service instance; and

each backup service instance includes means for notifying the client process that each backup service instance has become a new primary service instance; and

the new primary service instance includes means for receiving client requests and, for client requests that have already been committed by the new primary service instance when the new primary service instance was a backup service instance, providing the committed response to the client request without again processing the request.

14. The service of claim 13, wherein:

the new primary service instance includes means for receiving client requests and, for client requests that have already been received but not committed by the new primary service instance when the new primary service instance was a backup service instance, allowing the new primary service instance to continue processing the requests.

15. The service of claim 7, wherein:

a sequence of associated client requests is a transaction;

each backup service instance includes means for communicating a single acknowledgement for the client requests of the transaction; and

each backup service instance includes means for committing the results of processing the client requests of the transaction as a single commitment.

16. The service of claim 14, wherein:

a sequence of associated client requests is a transaction;

each backup service instance includes means for communicating an acknowledgement for each client request of the transaction; and

- 5 each backup service instance includes means for committing the results of processing the client requests of the transaction as a single commitment.

17. The service of claim 16, wherein:

the transaction committing means of each backup service instance includes means for
10 receiving a commitment request from the primary service instance, wherein the transaction committing means commits the results of the transaction in response to the received commitment request.

18. The service of claim 1, further including:

- 15 means for creating an additional backup service instance, including means for replicating the logical view of the primary service instance to the additional backup service instance.

19. The service of claim 18, wherein:

the means for replicating the logical view of the primary service instance to the additional
20 backup service instance includes providing the determined logical requests to the additional backup service instance.

20. The service of claim 19, wherein:

the replicating means includes means for processing a request history table to ensure that logical view of the additional backup service instance is the same as the logical view of the primary service instance.

5

21. The service of claim 20, wherein:

the replicating means includes means for configuring the determining means to process client requests received while the request history table is being processed.

10

22. The service of claim 21, wherein the means for configuring the determining means to process client requests received while the request history table is being processed operates in response to a termination of the request history table being processed.

15

23. The service of claim 22, wherein the transitioning means includes means for configuring the primary service instance to temporarily stop accepting client requests.

20

24. The service of claim 23, wherein the means for configuring the replacement primary service instance to temporarily stop accepting client requests operates in response to a comparison of the logical view of the primary service instance to the logical view of the additional backup service instance.

25. A service operating on a computer system to process requests from client processes, comprising:

a primary service instance operating on the computer system, including means for receiving and processing client requests from the client processes;

5 at least one backup service instance operating on the computer system, including means for receiving and processing the client requests, wherein each backup service instance is logically equivalent to the primary service instance;

wherein the primary service instance includes

10 means for determining which external requests, when processed by the secondary service instance, are logical requests such that processing of the determined logical requests cause the external view of each backup service instance to change; and

means for communicating with each backup service instance to provide an indication of the determined logical requests to each backup service instance.

15 26. The service of claim 25, further comprising:

means for creating and operating an additional backup service instance, including means for replicating the logical view of the primary service instance to the additional backup service instance.

20 27. The service of claim 25, wherein:

the means for creating and operating an additional backup service instance includes means for creating and operating the additional backup service instance on computing equipment having the same configuration as equipment on which the primary service instance is operating.

28. The service of claim 25, wherein:

the means for creating and operating an additional backup service instance includes means for creating and operating the additional backup service instance on computing equipment
5 having a different configuration as equipment on which the primary service instance is operating.

29. The service of claim 26, wherein:

the means for replicating the logical view of the primary service instance to the additional
10 backup service instance includes means for providing the determined logical requests to the additional backup service instance.

30. The service of claim 29, wherein:

the replicating means includes means for processing a request history table to ensure that
15 logical view of the additional backup service instance is the same as the logical view of the primary service instance.

31. The service of claim 26, wherein:

the means for replicating the logical view of the primary service instance to the additional
20 backup service instance includes means for implementing the additional backup service instance to have a physical behavior different from the physical behavior of the primary service instance.

32. The service of claim 31, further comprising:

means for disabling the at least one backup service instance;

means for enabling and operating the at least one backup service instance with an implementation that has the same physical behavior as the additional backup service

5 instance; and

means for replicating the logical view of the primary service instance to the at least one backup service instance.

33. The service of claim 32, further comprising:

10 means for disabling the primary service instance; and

means for causing at least one of the at least one backup service instance and the additional backup service instance to become a primary service instance.

34. The service of claim 33, further comprising:

15 means for creating and operating a second additional backup service instance, including

means for replicating the logical view of the primary service instance to the second additional backup service instance.

35. The service of claim 26, wherein:

20 the means for replicating the logical view of the primary service instance to the additional backup service instance includes means for implementing the additional backup service instance to have a physical behavior the same as the physical behavior of the primary service instance.

36. The service of claim 35, further comprising:

means for disabling the at least one backup service instance;

means for enabling and operating the at least one backup service instance with an
5 implementation that has the same physical behavior as the additional backup service
instance; and

means for replicating the logical view of the primary service instance to the at least one
backup service instance.

10 37. The service of claim 36, further comprising:

means for disabling the primary service instance; and

means for causing at least one of the at least one backup service instance and the additional
backup service instance to become a primary service instance.

15 38. The service of claim 37, further comprising:

means for creating and operating a second additional backup service instance, including
means for replicating the logical view of the primary service instance to the second additional
backup service instance.

20 39. A method of operating a computer system to process requests from client processes,
comprising:

operating a primary service instance on the computer system, including receiving and
processing client requests, from the client processes;

operating at least one backup service instance on the computer system, including receiving and processing the client requests, wherein each backup service instance is logically equivalent to the primary service instance;

wherein the step of operating the primary service instance includes

5 determining which external requests, when processed by the secondary service instance, are logical requests such that processing of the determined logical requests cause the external view of each backup service instance to change; and

 communicating with each backup service instance to provide an indication of the determined logical requests to each backup service instance.

10 40. The method of claim 39, wherein each backup service instance has a physical behavior different from the physical behavior of the primary service instance.

15 41. The method of claim 39, wherein a failure domain of the primary service instance is substantially independent of a failure domain of each backup service instance.

 42. The method of claim 39, wherein the step of providing the determined logical requests for the primary service instance includes providing the determined logical requests to each backup service instance via a network protocol.

20 43. The method of claim 39, wherein the step of operating each backup service instance includes persistently storing the determined logical requests.

44. The method of claim 39, wherein the step of operating each backup service instance includes communicating, to the primary service instance, acknowledgement of the determined logical requests.

45. The method of claim 39, wherein the step of operating the primary service instance includes committing the determined logical requests in response to the acknowledgements of the determined logical requests.

46. The method of claim 45, wherein the step of operating each backup service instance includes committing the results of processing the determined logical requests.

47. The method of claim 46, further comprising:
ensuring that, for a sequence of determined logical requests, each backup service instance commits the results of processing the determined logical requests in the same order as the primary service instance commits the results.

48. The method of claim 47, wherein:
step order ensuring step includes maintaining a table of determined logical requests; and
the step of operating each backup service instance includes processing the table of determined logical requests to determine an order of committing results of processing the determined logical requests.

49. The method of claim 48, wherein:

the step of maintaining a table of determined logical requests includes associating sequence values with the determined logical requests; and

the step of processing the table includes processing the sequence values.

5

50. The method of claim 39, wherein:

the step of operating the primary service instance includes receiving an indication of a path failure between the primary service instance and a client process that occurs after a result of processing a particular client request is committed and before the client process receives the result from the service; and

the step of operating the primary service instance includes receiving the particular client request again and providing the response to the client request without again processing the client request.

51. The method of claim 39, further comprising:

the step of operating each backup service instance includes receiving an indication of a failure of the primary service instance; and

the step of operating each backup service instance includes notifying the client process that each backup service instance has become a new primary service instance; and

the step of operating the new primary service instance includes receiving client requests and, for client requests that have already been committed by the new primary service instance when the new primary service instance was a backup service instance, providing the committed response to the client request without again processing the request.

52. The method of claim 51, wherein:

the step of operating the new primary service instance includes receiving client requests and, for client requests that have already been received but not committed by the new primary
5 service instance when the new primary service instance was a backup service instance, allowing the new primary service instance to continue processing the requests.

53. The method of claim 45, wherein:

a sequence of associated client requests is a transaction;

10 the step of operating each backup service instance includes communicating a single acknowledgement for the client requests of the transaction; and

the step of operating each backup service instance includes committing the results of processing the client requests of the transaction as a single commitment.

15 54. The method of claim 52, wherein:

a sequence of associated client requests is a transaction;

the step of operating each backup service instance includes communicating an acknowledgement for each client request of the transaction; and

the step of operating each backup service instance includes committing the results of
20 processing the client requests of the transaction as a single commitment.

55. The method of claim 59, wherein:

the step of committing a transaction for each backup service instance includes receiving a commitment request from the primary service instance, wherein the transaction committing step includes committing the results of the transaction in response to the received commitment request.

56. The method of claim 39, further including:

creating an additional backup service instance, including replicating the logical view of the primary service instance to the additional backup service instance.

57. The method of claim 56, wherein:

the step of replicating the logical view of the primary service instance to the additional backup service instance includes providing the determined logical requests to the additional backup service instance.

58. The method of claim 57, wherein:

the step of replicating the logical view of the primary service instance includes processing a request history table to ensure that logical view of the additional backup service instance is the same as the logical view of the primary service instance.

59. The method of claim 58, wherein:

the step of replicating the logical view of the primary service instance includes configuring the determining step to process client requests received while the request history table is being processed.

60. The method of claim 59, wherein the step of configuring the determining step to process client requests received while the request history table is being processed is performed in response to a termination of the request history table being processed.

5

61. The method of claim 60, wherein the transitioning step includes configuring the replacement primary service instance to temporarily stop accepting client requests.

10

62. The method of claim 61, wherein the step of configuring the replacement primary service instance to temporarily stop accepting client requests is performed in response to a comparison of the logical view of the primary service instance to the logical view of the additional backup service instance.

15

63. A method of operating a service on a computer system to process requests from client processes, comprising:

operating a primary service instance on the computer system, including receiving and processing client requests from the client processes;

20

operating at least one backup service instance on the computer system, including receiving and processing the client requests, wherein each backup service instance is logically equivalent to the primary service instance;

wherein the step of operating the primary service instance includes

determining which external requests, when processed by the secondary service instance, are logical requests such that processing of the determined logical requests cause

the external view of each backup service instance to change; and

communicating with each backup service instance to provide an indication of the determined logical requests to each backup service instance.

5 64. The method of claim 63; further comprising:

creating and operating an additional backup service instance, including replicating the logical view of the primary service instance to the additional backup service instance.

65. The method of claim 63, wherein:

10 the step of creating and operating an additional backup service instance includes creating and operating the additional backup service instance on computing equipment having the same configuration as equipment on which the primary service instance is operating.

66. The method of claim 63, wherein:

15 the step of creating and operating an additional backup service instance includes creating and operating the additional backup service instance on computing equipment having a different configuration as equipment on which the primary service instance is operating.

67. The method of claim 64, wherein:

20 the step of replicating the logical view of the primary service instance to the additional backup service instance includes providing the determined logical requests to the additional backup service instance.

68. The method of claim 64, wherein:

the replicating step includes processing a request history table to ensure that logical view of the additional backup service instance is the same as the logical view of the primary service instance.

5

69. The method of claim 64, wherein:

the step of replicating the logical view of the primary service instance to the additional backup service instance includes implementing the additional backup service instance to have a physical behavior different from the physical behavior of the primary service instance.

10

70. The method of claim 69, further comprising:

disabling the at least one backup service instance;

enabling and operating the at least one backup service instance with an implementation that has the same physical behavior as the additional backup service instance; and

15 replicating the logical view of the primary service instance to the at least one backup service instance.

71. The method of claim 70, further comprising:

disabling the primary service instance; and

20 causing at least one of the at least one backup service instance and the additional backup service instance to become a primary service instance.

72. The method of claim 71, further comprising:

creating and operating a second additional backup service instance, including replicating the logical view of the primary service instance to the second additional backup service instance.

5 73. The method of claim 64, wherein:

the step of replicating the logical view of the primary service instance to the additional backup service instance includes implementing the additional backup service instance to have a physical behavior the same as the physical behavior of the primary service instance.

10 74. The method of claim 73, further comprising:

disabling the at least one backup service instance;

enabling and operating the at least one backup service instance with an implementation that has the same physical behavior as the additional backup service instance; and

15 replicating the logical view of the primary service instance to the at least one backup service instance.

75. The method of claim 74, further comprising:

disabling the primary service instance; and

20 causing at least one of the at least one backup service instance and the additional backup service instance to become a primary service instance.

76. The method of claim 75, further comprising:
creating and operating a second additional backup service instance, including replicating the logical view of the primary service instance to the second additional backup service instance.

5 77. A program storage device readable by a machine tangibly embodying a program of instructions executable by the machine to perform method steps for operating a computer system to process requests from client processes, said method steps comprising:

operating a primary service instance on the computer system, including receiving and processing client requests, from the client processes;

10 operating at least one backup service instance on the computer system, including receiving and processing the client requests, wherein each backup service instance is logically equivalent to the primary service instance;

wherein the step of operating the primary service instance includes

15 determining which external requests, when processed by the secondary service instance, are logical requests such that processing of the determined logical requests cause the external view of each backup service instance to change; and

communicating with each backup service instance to provide an indication of the determined logical requests to each backup service instance.

20 78. The program storage device of claim 77, wherein each backup service instance has a physical behavior different from the physical behavior of the primary service instance.

79. The program storage device of claim 77, wherein a failure domain of the primary service instance is substantially independent of a failure domain of each backup service instance.

5 80. The program storage device of claim 77, wherein the step of providing the determined logical requests for the primary service instance includes providing the determined logical requests to each backup service instance via a network protocol.

10 81. The program storage device of claim 77, wherein the step of operating each backup service instance includes persistently storing the determined logical requests.

15 82. The program storage device of claim 77, wherein the step of operating each backup service instance includes communicating, to the primary service instance, acknowledgement of the determined logical requests.

83. The program storage device of claim 82, wherein the step of operating the primary service instance includes committing the determined logical requests in response to the acknowledgements of the determined logical requests.

20 84. The program storage device of claim 83, wherein the step of operating each backup service instance includes committing the results of processing the determined logical requests.

85. The program storage device of claim 84, further comprising:
ensuring that, for a sequence of determined logical requests, each backup service instance commits the results of processing the determined logical requests in the same order as the primary service instance commits the results.

5

86. The program storage device of claim 85, wherein:
step order ensuring step includes maintaining a table of determined logical requests; and
the step of operating each backup service instance includes processing the table of determined logical requests to determine an order of committing results of processing the determined logical requests.

10

87. The program storage device of claim 86, wherein:
the step of maintaining a table of determined logical requests includes associating sequence values with the determined logical requests; and
the step of processing the table includes processing the sequence values.

15

88. The program storage device of claim 77, wherein:
the step of operating the primary service instance includes receiving an indication of a path failure between the primary service instance and a client process that occurs after a result of processing a particular client request is committed and before the client process receives the result from the service; and
the step of operating the primary service instance includes receiving the particular client request again and providing the response to the client request without again processing the client request.

20

89. The program storage device of claim 77, further comprising:

the step of operating each backup service instance includes receiving an indication of a failure of the primary service instance; and

5 the step of operating each backup service instance includes notifying the client process that each backup service instance has become a new primary service instance; and

the step of operating the new primary service instance includes receiving client requests and, for client requests that have already been committed by the new primary service instance when the new primary service instance was a backup service instance, providing the
10 committed response to the client request without again processing the request.

90. The program storage device of claim 89, wherein:

the step of operating the new primary service instance includes receiving client requests and, for client requests that have already been received but not committed by the new primary
15 service instance when the new primary service instance was a backup service instance, allowing the new primary service instance to continue processing the requests.

91. The program storage device of claim 83, wherein:

a sequence of associated client requests is a transaction;

20 the step of operating each backup service instance includes communicating a single acknowledgement for the client requests of the transaction; and

the step of operating each backup service instance includes committing the results of processing the client requests of the transaction as a single commitment.

92. The program storage device of claim 89, wherein:

a sequence of associated client requests is a transaction;

the step of operating each backup service instance includes communicating an

5 acknowledgement for each client request of the transaction; and

the step of operating each backup service instance includes committing the results of processing the client requests of the transaction as a single commitment.

93. The program storage device of claim 92, wherein:

10 the step of committing a transaction for each backup service instance includes receiving a commitment request from the primary service instance, wherein the transaction committing step includes committing the results of the transaction in response to the received commitment request.

15 94. The program storage device of claim 77, further including:

creating an additional backup service instance, including replicating the logical view of the primary service instance to the additional backup service instance.

95. The program storage device of claim 94, wherein:

20 the step of replicating the logical view of the primary service instance to the additional backup service instance includes providing the determined logical requests to the additional backup service instance.

96. The program storage device of claim 95, wherein:

the step of replicating the logical view of the primary service instance includes processing a request history table to ensure that logical view of the additional backup service instance is the same as the logical view of the primary service instance.

5

97. The program storage device of claim 96, wherein:

the step of replicating the logical view of the primary service instance includes configuring the determining step to process client requests received while the request history table is being processed.

10

98. The program storage device of claim 97, wherein the step of configuring the determining step to process client requests received while the request history table is being processed is performed in response to a termination of the request history table being processed.

15

99. The program storage device of claim 98, wherein the transitioning step includes configuring the replacement primary service instance to temporarily stop accepting client requests.

20

100. The program storage device of claim 99, wherein the step of configuring the replacement primary service instance to temporarily stop accepting client requests is performed in response to a comparison of the logical view of the primary service instance to the logical view of the additional backup service instance.

101. A program storage device readable by a machine tangibly embodying a program of instructions executable by the machine to perform method steps for operating a computer system to process requests from client processes, said method steps comprising:

5 operating a primary service instance on the computer system, including receiving and processing client requests from the client processes;

operating at least one backup service instance on the computer system, including receiving and processing the client requests, wherein each backup service instance is logically equivalent to the primary service instance;

10 wherein the step of operating the primary service instance includes

determining which external requests, when processed by the secondary service instance, are logical requests such that processing of the determined logical requests cause the external view of each backup service instance to change; and

15 communicating with each backup service instance to provide an indication of the determined logical requests to each backup service instance.

102. The program storage device of claim 101, further comprising:

creating and operating an additional backup service instance, including replicating the logical view of the primary service instance to the additional backup service instance.

20 103. The program storage device of claim 101, wherein:

the step of creating and operating an additional backup service instance includes creating and operating the additional backup service instance on computing equipment having the same configuration as equipment on which the primary service instance is operating.

104. The program storage device of claim 101, wherein:

the step of creating and operating an additional backup service instance includes creating and operating the additional backup service instance on computing equipment having a different configuration as equipment on which the primary service instance is operating.

105. The program storage device of claim 102, wherein:

the step of replicating the logical view of the primary service instance to the additional backup service instance includes providing the determined logical requests to the additional backup service instance.

106. The program storage device of claim 105, wherein:

the replicating step includes processing a request history table to ensure that logical view of the additional backup service instance is the same as the logical view of the primary service instance.

107. The program storage device of claim 102, wherein:

the step of replicating the logical view of the primary service instance to the additional backup service instance includes implementing the additional backup service instance to have a physical behavior different from the physical behavior of the primary service instance.

108. The program storage device of claim 107, further comprising:

disabling the at least one backup service instance;

enabling and operating the at least one backup service instance with an implementation that has the same physical behavior as the additional backup service instance; and

5 replicating the logical view of the primary service instance to the at least one backup service instance.

109. The program storage device of claim 108, further comprising:

disabling the primary service instance; and

10 causing at least one of the at least one backup service instance and the additional backup service instance to become a primary service instance.

110. The program storage device of claim 109, further comprising:

creating and operating a second additional backup service instance, including replicating the

15 logical view of the primary service instance to the second additional backup service instance.

111. The program storage device of claim 102, wherein:

the step of replicating the logical view of the primary service instance to the additional backup service instance includes implementing the additional backup service instance to

20 have a physical behavior the same as the physical behavior of the primary service instance.

112. The program storage device of claim 111, further comprising:

disabling the at least one backup service instance;

enabling and operating the at least one backup service instance with an implementation that has the same physical behavior as the additional backup service instance; and

5 replicating the logical view of the primary service instance to the at least one backup service instance.

113. The program storage device of claim 112, further comprising:

disabling the primary service instance; and

10 causing at least one of the at least one backup service instance and the additional backup service instance to become a primary service instance.

114. The program storage device of claim 113, further comprising:

creating and operating a second additional backup service instance, including replicating the

15 logical view of the primary service instance to the second additional backup service instance.